

# MECHANISM FOR NESTED EXPANSION OF DATA COLLECTION

## Technical Field

The present invention relates to the field of determining information from multiple computers connected on a system. More specifically, the present invention relates to the gathering of information utilized on a system based on a Common Information Modeling (CIM) standard.

## Background

Commonly, pluralities of computers, databases, printers and other computing or computing related devices are often established in clusters, as members or elements of a network, or even as loosely defined connections (which may be permanent or temporary) of computers and/or computing devices. Such configurations shall hereafter collectively be identified as a Cluster. Similarly, those computing devices establishing the connections between and with other such devices shall be defined as Nodes. Further, it is commonly appreciated that various virtual and/or real devices and components are often utilized in and/or accessed by a computing device (for example, printers, hard drives, monitors, software applications and others). For purposes of the present discussion, such devices are herein identified as Resources. Such Resources are generally accessible and/or utilized via other systems and devices, and/or users of such systems and devices via Nodes on a Cluster. Such users (which may be human or automated) are herein collectively defined as Clients. As such, it is commonly appreciated that Clients, via Nodes on Clusters, often utilize and need information pertaining to all the Resources or specific types or instances of Resources accessible to the Client.

Since Clients often desire to know what type of Resources are accessible via a given Cluster, various software devices have been developed that enable Clients to obtain information pertaining to a Resource, a Node, a Cluster or a combination thereof. Such devices are often called Object Managers (OM). More specifically, standardized OMs have been developed, including for example, the Common Information Modeling standard (CIM). The CIM standard for OMs was created by a consortium of industry companies and can be found at <http://www.dmtf.org>. It is commonly appreciated that CIM and WBEM are extremely similar, as such CIM is herein defined to also include WBEM.

The CIM standard outlines the basic architecture for a CIM compliant OM (CIMOM) but does not enforce any implementation details. As is commonly appreciated,

1 CIM enables application developers to code their Client's applications so that they may  
2 interact with the CIMOM using a predefined, platform independent, set of Application  
3 Program Interfaces (APIs). This basic CIM architecture enables a CIMOM to startup and  
4 load any CIM compliant Resource via a Provider, wherein a Provider is a piece of  
5 software code that will call whatever routines are needed to gather data in order to  
6 provide the necessary objects in a schema. Those skilled in the art appreciate that a  
7 schema is an abstraction of something that exists in the real world. As such, in CIM, a  
8 schema is a group of classes (collections or sets of objects) that have only one owner. For  
9 example, a collection of disc drives might have an owner identified as a specific database,  
10 i.e., a Resource, that is connected to the Cluster and is accessible via the associated  
11 Provider.

12 Each Provider is suitably configured to extract information from the various  
13 objects with which it may be associated. For example, in a multi-disc or tape database, a  
14 database Provider is configured to obtain necessary information from each object in the  
15 database. Such information might include available memory, file allocations, and other  
16 information. There are often hundreds if not thousands of Providers on a given Node  
17 and/or Cluster. Provider writers commonly come up with a schema of what they are  
18 going to provide, and then code their respective Providers to provide such objects upon  
19 request. The CIM standard defines those APIs utilized to interact and extract information  
20 from a Provider. These standards are utilized by Provider writers.

21 As such, it is to be appreciated that a pre-defined methodology exists for  
22 identifying information about a Resource, a group of Resources, or even a sub-set of a  
23 Resource to an OM under the CIM standard. Under these standards, each CIMOM is  
24 configured such that it can determine which Provider is providing the desired information  
25 (for example, the amount of available memory in a database). Further, the CIMOM is  
26 configured to "ask" for and receive such information from a Provider. However, the CIM  
27 standard does not specify how information from Providers is to be obtained. Since  
28 Cluster managers often desire Cluster-wide configuration and Resource information, an  
29 approach is needed to determine how Resources are configured on a Cluster using  
30 Providers and requests from CIMOMs.

31 One common methodology by which a Client CIM can obtain Cluster-wide  
32 information concerning Resources is shown in Figure 1. As shown, this prior art process  
33 basically entails having the Client CIM contact every Node on the Cluster. As can be  
34 readily appreciated, this approach requires numerous connections between the requesting

1 Client CIM 102 and a plurality of Node CIMOMs (110, 112, and 114) on Nodes A, B and  
2 C (104, 106 and 108, respectively). In turn, each of the Node CIMOMs, may have  
3 multiple Providers connected thereto (for example, Providers A and B, i.e., 116, 118, 120,  
4 122, 124, and 126). For even a simple three Node example, where each Node possesses  
5 two Providers (as shown in Figure 1), the request from the Client CIM 102 over the  
6 communications network 128 entails requests to three separate Node CIMOMs, which  
7 may end up requesting information from six Providers. In an extremely large system,  
8 wherein hundreds if not thousands of Nodes may exist, this approach can be extremely  
9 time and Resource consuming. This may be especially true when the Client CIM 102 can  
10 only contact one Node at a time and, thus, is forced to contact each Node on a Cluster in a  
11 serial manner. Further, since the network connection 128, in certain embodiments, may  
12 not be secure, this process may also entail numerous authorizations and authentication  
13 steps being performed with each request by the Client CIM 102 to a Node CIMOM (for  
14 example, one provided on Node A 104, Node B 106 or Node C 108). These  
15 authorizations may also create delays in the processing of the request and the operating  
16 efficiency of the Cluster as a whole. Therefore, the approach shown in Figure 1 is often  
17 tedious and inefficient. This approach is often considered as operating at too high a level  
18 of abstraction and is therefore undesirable because of the burdens and inefficiencies it  
19 places upon the Client CIM 102 and the Cluster in general.

20 Another prior art approach often utilized to obtain information from Providers by  
21 a Client CIM 202 is shown in Figure 2. This approach again begins with the Client CIM  
22 202 establishing a connection via a communications network 228 with another Node (for  
23 example, Node B 206) on the Cluster. However, unlike the example shown in Figure 1,  
24 for this embodiment only a single connection is established between the Client CIM 202  
25 and a single Node (for example, Node B 206). In short, this approach does not require the  
26 Client CIM 202 to contact each and every Node on the Cluster. Instead, this approach  
27 utilizes Providers that are multi-node aware. One example of a multi-node aware  
28 Provider is provided in United States Patent Application Serial Number concurrently filed  
29 herewith, on August 31, 2001, by Jim Curtis, et al. and entitled "An Application  
30 Container that Allows Concurrent Execution on Multiple Nodes in a Cluster", the  
31 contents of which are herein incorporated, in their entirety, by reference.

32 As shown in Figure 2, this second prior art method requires each Provider to  
33 include a data daemon (for example, data daemons 230, 232, 234, 236, 238 and 240).  
34 Each data daemon contains the appropriate software codes such that when a request from

1 a single Provider (for example, Provider A 220 of Node B 206) is communicated to the  
2 various other Nodes on the Cluster, each corresponding data daemon (for example, data  
3 daemon A 230 on Node A, or data daemon A 234 on Node B 206, or data daemon A 238  
4 on Node C) provides the requested information, when available. In this prior art  
5 embodiment, every Provider, and thus every data daemon, must be Cluster-aware.

Noticeably, even the requesting Node (Node B 206) contains data daemons 234  
and 236 from which configuration information may be needed. Further, this approach  
requires each Provider to be capable of contacting other Providers on the Cluster. Since  
certain applications may not provide Cluster-aware Providers, this approach may be of  
limited applicability. Further, instead of having the Client contact every Node on the  
Cluster, this approach merely reduces the decision/contacting processes to the Provider  
level. Further, this approach requires every Provider to be capable of multiplexing  
requests to every other Node on the Cluster. As such, this approach reduces the level of  
abstraction to too low a level and requires every Provider to possess capabilities  
previously provided by the requesting Client CIM, as shown in Figure 1. Therefore, a  
new system and methodology for obtaining information from a Provider is needed which  
does not place the burden for determining information about Resources on a Cluster at  
either the Client CIM (i.e., too high of a level of abstraction) or at the Provider level (i.e.,  
too low of a level of abstraction).

## Summary

The present invention provides a methodology for collecting information from a plurality of computers, on a Cluster, by utilizing a Multiplex Provider (MP). The MP enables single-system CIM compliant OM Providers to be used, without modification, to provide multi-node data from a Cluster(s). By utilizing the MP, a Client can connect to a single Node on the Cluster and obtain information on several different peer Nodes, via a Node CIMOM and an MP, without having to establish a direct connection to each Node on the Cluster and without requiring every Provider to be Cluster-aware.

Further, the MP avoids putting the burden of multiplexing on either the Client or the Provider writer. The Client requests multi-node data through a scoping hint that is passed along with the request to a Node CIMOM. This scoping hint identifies those Nodes from which the Client desires to receive the requested information. The Node CIMOM suitably inspects the request in order to determine which classes/schemas of objects are being requested and which Providers provide those classes/schemas. With the MP mechanism, the Node CIMOM can also determine which classes/schemas should be

1 multiplexed, and utilizes the scoping hints to decide with which Nodes to open a  
2 connection. In this manner, the MP effectively acts as a proxy for the Node CIMOM.  
3 The MP takes on the role of a Client CIM (as discussed previously with reference to  
4 Figure 1) by requesting information from other Nodes on the Cluster, while also  
5 removing such tasks to a lower level of abstraction such that the Client CIM is not  
6 directly contacting every Node on the Cluster with requests for information.  
7 Additionally, in its requests to the other Nodes, the MP does not specify a scoping hint so  
8 as to not trigger the MPs on the other Nodes to intercept the request. As such, by not  
9 providing the scoping hints, a single set of data for the Cluster, from the desired Nodes,  
10 may be provided to the Client by the MP.

11       In addition to streamlining the process for requesting data and alleviating the  
12 burden from the Client and the Providers, the MP also provides data filtering capabilities.  
13 Specifically, once the MP receives responses from the other Nodes, the MP gathers the  
14 data, eliminates any redundant data, and provides the information (preferably in a tabular  
15 form) back to the Client CIM which appropriately communicates the information back to  
16 the Client. Further, the MP may also be configured to manage other related processing  
17 steps, such as authorizations and authentications to remote Nodes.

18       As such, the present invention provides an improved methodology for obtaining  
19 information from Resources on a Cluster. The methodology is preferably configured to  
20 be implemented on a CIM environment, however, it may also be accomplished in other  
21 environments, including, but not limited to, WBEM. The features and functions of the  
22 present invention are further described in the drawing figures, the detailed description and  
23 the claims.

#### 24                   **Description of the Drawings**

25       Figure 1 provides an illustration of a prior art method commonly utilized to obtain  
26 information about Resources on a Cluster upon a Client's request.

27       Figure 2 provides an illustration of a second prior art method commonly utilized  
28 to obtain information about Resources on a Cluster upon a Client's request.

29       Figure 3 provides an illustration of the process of the present invention wherein an  
30 MP is utilized to obtain information about Resources on a Cluster upon a Client's request.

31       Figures 4A and 4B provide a flow diagram of the process by which the MP  
32 obtains information from the various Nodes on a Cluster.

## Detailed Description

2 As mentioned previously, the present invention utilizes an MP to obtain  
3 information on Resources, from Nodes on a Cluster, upon a Client's request. Unlike the  
4 prior art methods, the present invention does not require the Client to contact every Node  
5 on the network in order to obtain information about Resources available via the Cluster.  
6 Thus, the present invention alleviates many of the data connectivity and data obtaining  
7 tasks from the Client and shifts such tasks down to the MP, where they can be more easily  
8 implemented and with less Client to multiple Node communications being needed. It is  
9 to be appreciated, that by utilizing this approach the data traffic on the network 328 (see  
10 Figure 3) may be significantly reduced since every Client is not continually contacting  
11 every Node on the Cluster whenever Resource information is desired.

12        Additionally, unlike the prior art method shown in Figure 2, the present invention  
13      does not require every Provider to be Cluster-aware. Instead, the MP provides the needed  
14      Cluster-awareness. Since an MP is preferably loaded onto each Node, Providers that are  
15      not Cluster-aware may be suitably added to a Node while information about such  
16      Provider may be obtained from other Providers or Clients by the MP. As such, the  
17      present invention greatly simplifies the adding of Providers to Nodes, in that such  
18      Providers do not need to be or become Cluster-aware. The Cluster-awareness is  
19      preferably provided by the MP associated with each Node on the Cluster. Further, the  
20      present invention expands the possible number of Providers that may be utilized on a  
21      Cluster, because all that is required, in order to be “known” on the Cluster, is the ability to  
22      interface with the MP without requiring any capabilities to “know” of other Providers on  
23      the same or other Nodes.

Further, the MP approach of the present invention also provides enhanced update capabilities. Instead of having to update every Provider with Cluster information, as provided for with the embodiments shown in Figure 2, the present invention enables Cluster managers to only update the MP associated with a Node or multiple Nodes on the Cluster. It is readily appreciated that updating the MP on each Node is less extensive of an operation than having to update every Provider on a Cluster. Similarly, economies of scale realizations may also be experienced with regard to verifications, authentications, authorizations, and other processes which may need updating or augmenting. The present invention generally removes the target of such updates from the individual Providers and places it on the MP, thereby, further incurring savings in time and resources needed to maintain and/or operate a Cluster.

1 Referring now to Figure 3, one embodiment of an implementation of the process  
2 of the present invention is depicted. As shown, this embodiment utilizes a Client CIM  
3 302 that has established a connection, via a network 328, with a Node (in this case Node  
4 B 306) on a Cluster 300. The Node 306 suitably includes a CIMOM 312 that includes an  
5 MP 334. The MP 334 is suitably connected via data links 330 and 332 to the various  
6 other Nodes (Node A 304 and Node B 308) on the Cluster 300. It is to be appreciated  
7 that such data links 330 and 332 may be same links utilized by the Node 306 to  
8 communicate with its peer Nodes or any other links. Further, the data links 330 and 332  
9 may also be separate links, specifically configured for communications between MPs and  
10 the other Nodes on the Cluster and/or the CIMOMs associated with such other Nodes.  
11 Thus, any link between the MP and the additional CIMOM Nodes on the Cluster may be  
12 utilized by the present invention.

13 Further, the MP 334 communicates directly with the CIMOM on each Node (310  
14 and 314, respectively), thereby alleviating the burden of such communications from the  
15 Client CIM 302 and/or the Providers (316, 318, 320, 322, 324, and 326). It is to be  
16 appreciated that each Node, in the preferred embodiment, includes an MP 334. As such,  
17 a Client does not need to connect to a specific Node in order to receive Resource related  
18 information that is obtained Cluster-wide, or for even a limited sub-set of a Cluster.  
19 However, the present invention may be configured in alternative embodiments, as  
20 desired, such that only specific Nodes contain the MP and provide the capabilities of  
21 requesting and obtaining information from the various Providers on the various Nodes, or  
22 specific Nodes, of a Cluster.

23 Referring now to Figures 4A and 4B (with reference to Figure 3), one  
24 embodiment of a process, by which the present invention utilizes an MP to obtain  
25 information on Resources associated with a Cluster, is shown. The process begins with a  
26 Client connecting to a Node 306 via a Client CIM 302 (Block 402). It is to be  
27 appreciated, that each Client is suitably configured with a CIM compatible interface that  
28 allows the Client to establish communications with a particular Node on the Cluster (for  
29 example, Node B 306, as shown in Figure 3). When connecting to the Node 306, the  
30 Client CIM 302 suitably performs those commonly known in the art tasks such as  
31 authenticating, verifying connectivity, and other such tasks.

32 Once the connection between the Client (via the Client CIM 302) and the Node  
33 306 is established, the process continues with the Client sending a query to the CIMOM  
34 312 on the Node 306 (Block 404). It is to be appreciated, that depending upon the

1 specific implementation of the process of the present invention, that a request may have  
2 to be specified in a particular format in order to be valid. One such format includes  
3 HTTP, but, the present invention is not so limited. Further, the query suitably contains a  
4 request(s) for information about either the Resources connected to the Node 306 and/or  
5 Resources connected to other Nodes (for example, Node A 304 and Node C 308) on the  
6 Cluster 300. At this point, the CIMOM 312 determines whether the query from the Client  
7 is a request for Resource information or a request for other information or services (for  
8 example, a request to lock-up a Resource) (Block 406). If the query is for non-Resource  
9 information or services provided by the Node 306, the CIMOM 312 processes the request  
10 using established procedures that are beyond the scope of the present invention (Block  
11 407).

12 When the query is for Resource information, the CIMOM 312 suitably passes the  
13 query to the MP 334 (Block 408). The MP 334 then determines the scope of the query  
14 (Block 410). It is to be appreciated that a query from a Client may request information on  
15 every Resource on a Cluster, a limited type of Resources (for example, only information  
16 on disc drives), information on limited Resources as identified by specific Nodes, and/or  
17 based upon any other criteria. In order to determine from which Nodes a Client may  
18 desire to request information, the present invention enables a Client to request discovery  
19 on the Cluster 300 (Block 412).

20 When discovery is requested, the CIMOM 312, via the MP 334, suitably  
21 communicates a broadcast probe across the Cluster 300 (Block 414). The broadcast  
22 probe may be a general probe wherein, for example, every CIMOM Node (for example,  
23 Node A 304 and Node C 308, in Figure 3) identifies itself to the requesting CIMOM  
24 Node (Node B 306). Alternatively, any level of specificity may be utilized in requesting  
25 CIMOM Nodes and/or Providers (which, in certain embodiments, may even include those  
26 Providers connected to the requesting CIMOM Node) to identify their existence to the  
27 requesting CIMOM Node. For example, a broadcast probe might seek only those  
28 Providers that are connected to data storage Resources. Similarly, other broadcast probes  
29 might seek processors capable of performing specific operations. Thus, the broadcast  
30 probe may be for any topic and sent to any desired distribution of Nodes, Clusters,  
31 Resources, and/or Providers.

32 The requesting CIMOM Node (for this example, Node B 306) suitably receives  
33 replies from the responding CIMOM Nodes (e.g., Node A 304 and Node C 308) and  
34 compiles a table that includes a listing of such responses (Block 416). The CIMOM 312

1 then supplies this information, preferably, but not necessarily, in tabular form, to the  
2 Client CIM 302 (Block 418). The Client CIM then suitably specifies a scope for the  
3 query (Block 420) based upon the listings provided by the CIMOM 312. The scope may  
4 specify specific, or all, Nodes, Resources, Providers, real or virtual elements, sub-  
5 components and/or aspects that are to receive the query. Further, the scope may also  
6 specify that certain CIMOM Nodes are to receive certain aspects of the query, while other  
7 CIMOM Nodes receive other aspects of the query. This may be desirable when a query is  
8 a compound query that seeks identification of multiple Resources and/or Providers that  
9 may or may not provide unique or different functions and/or services.

10 Further, it is to be appreciated that for certain queries, wherein large amounts of  
11 information may be requested from multiple CIMOM Nodes and/or Providers, sub-sets of  
12 queries might be desired and/or multiple queries might be utilized. The present invention  
13 suitably supports any scope and may be configured, as desired by specific applications, to  
14 support queries which utilize compounded scopes or otherwise. The present invention  
15 provides Clients, CIMOMs and/or MPs with the capabilities to submit queries of any  
16 desired scope. In short, the scope effectively allows the Clients, CIMOM and/or MPs to  
17 specify, in advance, from whom they are going to request information.

18 Additionally, the present invention supports the concept of pre-defined/default  
19 scopes. Such a scope might include, for example, all of the local CIMOM Nodes on a  
20 Cluster, or all of the Providers connected to certain CIMOM Nodes, and the like. These  
21 pre-defined scopes may be utilized at any time and may be specified by a Client as  
22 necessary and/or desired. It is to be appreciated that scopes may be suitably stored in  
23 almost any available data storage device known in the art and retrieved accordingly, as  
24 desired and/or necessary, by the Client CIMOM and/or a CIMOM Node.

25 Referring again to Figures 4A and 4B, the process continues after Block 420 with  
26 the CIMOM 312 communicating the request and the scope of the request to the MP 334  
27 (Block 422). The MP 334 then requests classes of objects from the CIMOMs on the  
28 respective CIMOM Nodes (in this example, Nodes A 304 and B 308) and from the  
29 Providers associated with such CIMOM Nodes (i.e., 316, 318, 320, 322, 324 and 326)  
30 (Block 424). The request communicated from the MP 334 to the CIMOM Nodes is  
31 preferably communicated with a local scope. This local scope effectively instructs the  
32 MP on each of the CIMOM Nodes to not conduct a Cluster-wide search and, instead, to  
33 conduct a search only of those Providers with which each CIMOM Node is associated.  
34 For example, a local scope request communicated to CIMOM Node A 304 would instruct

1 the CIMOM Node to only communicate the request to Provider A 316 and/or Provider B  
2 318, i.e., the Providers directly connected or associated with CIMOM Node A 304. Such  
3 a local scope request would not obtain information for Providers connected to CIMOM  
4 Node C 308. If information was desired from CIMOM Node C, the MP 334 would need  
5 to communicate a local scope request directly to CIMOM Node C 308 in order to obtain  
6 information, for example, pertaining to Provider A 324. Commonly, this process is  
7 further supported by the knowledge of the particular classes/schemas of objects held by  
8 each CIMOM for its related Providers, such that each CIMOM Node does not spend  
9 unnecessary time contacting Providers incapable of providing the requested information.

10 More specifically, when a Provider is started on a Node, the CIMOM for that  
11 Node obtains from the Provider the types of classes/schemas held by the Provider. This  
12 information is then suitably stored by the CIMOM for later utilization. Further, a  
13 determination is made as to whether such classes are to be multiplexed, i.e., available to  
14 the MP. If so, then these classes are suitably stored in a multiplex table associated with  
15 the CIMOM Node. When the request comes from the MP on the requesting CIMOM  
16 Node, each CIMOM Node receiving the request utilizes the class/schema information  
17 previously extracted from the Providers and stored in their respective multiplex tables to  
18 determine to which Providers to further communicate the query. As such, for the present  
19 embodiment, the process is preferably utilized in conjunction with MP. But, the process  
20 may also be utilized with non-multiplexed Providers, provided the requesting CIMOM  
21 Node includes, and utilizes, an MP. Further, it is to be appreciated that non-multiplexed  
22 Providers may still be accessed by a system implementing the process of the present  
23 invention. However, such access is usually provided by having the MP directly connect  
24 to a Provider utilizing a process similar to one of the prior art processes previously  
25 described herein or a comparable process.

26 Referring again to Figures 4A and 4B, after the request is communicated by the  
27 MP 334 to each CIMOM Node on the Cluster (or a limited set thereof) with a local scope,  
28 the process continues with such CIMOM Nodes returning the desired information, if  
29 available, back to the MP 334 on the requesting CIMOM Node (i.e., Node B 306) (Block  
30 426). The MP 334 on the requesting CIMOM Node 306 then combines the various  
31 results received from the various CIMOM Nodes/Providers into tables which are suitably  
32 combined as desired and/or necessary (Block 428). The tables may be combined in any  
33 manner specified by the Client. Such manners might include showing all instances  
34 received on various tables and/or only showing unique entries. Other combinations

1 and/or filters may also be utilized to sort, organize and/or present requested and received  
2 information to a Client. Also, commonly known query languages and techniques can be  
3 utilized to accomplish any desired filtering of received data, as desired by the Client or  
4 Cluster manager.

5 Upon production of the tables by the requesting CIMOM, the process then  
6 continues with the Client CIM 302 receiving the completed table (or, in certain instances,  
7 portions thereof) and communicating such results to the Client (Block 430). At this point  
8 in the process, the Client has then received the requested information.

9 At this point, the process ends (Block 432) until another request is generated by  
10 the same or a different Client. As such, it is to be appreciated that the process flows  
11 depicted in Figure 4 may occur simultaneously on multiple nodes of a network, when  
12 such Nodes are suitably configured to operate in conjunction with an MP.

13 Therefore, the present invention has been described in the context of a process  
14 flow implemented on a system that utilizes CIMs, CIMOMs, MPs and Providers to  
15 communicate with, request and receive information about objects/schemas on a Cluster  
16 without requiring Cluster-aware Providers and/or Clients that are capable of  
17 communicating with every Node on a Cluster. Further, while the present invention has  
18 been described in the context of specific embodiments and process flows, it is to be  
19 appreciated that the present invention is not so limited and may encompass any systems  
20 and/or process flows which are within the scope of the present invention as specified by  
21 the following and any subsequently added claims.